

Data-Driven Discovery of Biophysical T Cell Receptor Cospecificity Rules

Andrew G. T. Pyo,¹ Yuta Nagano^{2,3}, Martina Milighetti^{2,4}, James Henderson^{2,5}, Curtis G. Callan, Jr.^{1,6}, Benny Chain,^{2,7} Ned S. Wingreen^{8,9} and Andreas Tiffeau-Mayer^{2,5}

¹Department of Physics, *Princeton University*, Princeton, New Jersey 08544, USA

²Division of Infection and Immunity, *University College London*, London, United Kingdom

³Division of Medicine, *University College London*, London, United Kingdom

⁴Cancer Institute, *University College London*, London, United Kingdom

⁵Institute for the Physics of Living Systems, *University College London*, London, United Kingdom

⁶Département de Physique, *École Normale Supérieure*, Paris, France

⁷Department of Computer Science, *University College London*, London, United Kingdom

⁸Lewis-Sigler Institute for Integrative Genomics, *Princeton University*, Princeton, New Jersey 08544, USA

⁹Department of Molecular Biology, *Princeton University*, Princeton, New Jersey 08544, USA



(Received 19 December 2024; accepted 18 June 2025; published 15 July 2025)

The biophysical interactions between the T cell receptor (TCR) and its ligands determine the specificity of the cellular immune response. However, the immense diversity of receptors and ligands has made it challenging to discover generalizable rules across the distinct binding affinity landscapes created by different ligands. Here, we present an optimization framework for discovering biophysical rules that predict whether TCRs share specificity to a ligand. Applying this framework to TCRs associated with a collection of SARS-CoV-2 peptides, we systematically characterize how cospecificity depends on the type and position of amino-acid differences between receptors. We also demonstrate that the inferred rules generalize to ligands highly dissimilar to any seen during training. Our analysis reveals that the matching of steric properties between substituted amino acids is more important for receptor cospecificity than the hydrophobic properties that prominently determine evolutionary substitutability. Our analysis also quantifies the substantial importance of positions not in direct contact with the peptide for specificity. These findings highlight the potential for data-driven approaches to uncover the molecular mechanisms underpinning the specificity of adaptive immune responses.

DOI: [10.1103/14j1-wrh5](https://doi.org/10.1103/14j1-wrh5)

I. INTRODUCTION

Cellular immunity relies on the specific recognition of target molecules by T cells, mediated by the binding of the T cell receptor (TCR) to specific peptide-major histocompatibility complexes (pMHCs) [1,2]. Given the pivotal role of T cells in the adaptive immune system, an ability to predict the specificity of a TCR *in silico* from its sequence would have many applications in disease diagnosis, surveillance, and treatment. Recent experimental advances [3–7] and the creation of databases cataloging known TCR-pMHC pairings [8–11] have thus propelled significant research efforts aimed at predicting TCR specificity using machine learning [12–20].

The premise behind using machine learning to predict pairings between the hypervariable receptors and their ligands is that there are learnable, generalizable biophysical rules that govern these interactions [3,4]. While machine learning models are able to accurately identify alternative binders to pMHCs for which good quality high-volume TCR data exist, an ability to generalize to pMHCs unseen during training has yet to be demonstrated in independent benchmarks [21].

In high-dimensional multiclass classification problems, a potentially easier path towards generalizable learning begins with learning similarity relationships between class members [22,23]. Once a suitable metric or representation tailored to the general nature of the classification problem is learned, it can then be adapted with little training data to new classes. However, to date the most popular sequence similarity metrics for TCRs are heuristic, and their parameters have not been learned from TCR sequence data. For instance, TCRdist [3] scores amino-acid similarity based on the BLOcks SUBstitution Matrix (BLOSUM). This matrix is based on substitution frequencies observed in general protein evolution [24], and it is unclear how well these describe the impact of amino-acid changes on TCR-pMHC binding. Additionally, existing structures suggest that factors beyond the identity of substituted amino acids, such as the position of the substitution, likely influence the impact of sequence changes [25].

To address these challenges, we propose a pseudolikelihood maximization approach to learn the biophysical rules that govern TCR cospecificity directly from data. Concretely, our aim is to learn a distance metric $d(\sigma, \sigma')$ that quantifies the probability that two TCRs, σ and σ' , bind to the same ligand (pMHC). To accomplish this, we expand on our ongoing efforts to quantify the statistics of TCR sequence similarity from available TCR-pMHC specificity data [26–28] by replacing predefined notions of sequence similarity with

learnable metrics. To learn this metric, we apply contrastive learning—a framework that we and others have previously used to train deep learning models for predicting TCR specificity [19,20,29,30]—to optimize sequence alignment weights to predict TCR cospecificity. We reason that the simplicity of these metrics might favor the inference of generalizable rules from limited data, which has remained a challenge in more flexible deep learning models using this supervised contrastive approach [20].

Our results show that cospecificity rules inferred from a large-scale dataset on SARS-CoV-2-specific TCRs [11] generalize with striking breadth. These rules accurately predict TCR cospecificity even for pMHCs that are highly dissimilar (differing by ≥ 6 amino acids) from the nearest peptide in the training set, demonstrating an extrapolation ability that has so far remained out of reach for sequence-based machine learning approaches [18,21]. Beyond the empirical findings, we provide a theoretical perspective on how contrastive learning is related to pseudolikelihood methods and intuition for why it might identify reproducible order within disordered receptor-ligand landscapes, where traditional inference methods often falter. Altogether, our findings position contrastive learning as a powerful and broadly applicable strategy for developing statistical theories of complex receptor-ligand maps from limited data.

II. TWO-POINT STATISTICS IDENTIFY ORDER IN RECEPTOR-LIGAND MAPS

To motivate our framework, it is useful to review the statistical inference problem posed by the TCR-pMHC binding prediction. The ultimate objective is to infer the probability distribution $P_\pi(\sigma)$ of observing the TCR sequence σ when sampling from the pool of T cells specific to pMHC π . This distribution can be fitted from known experimentally sampled TCR binders to the pMHC. However, there are only a few pMHCs with substantial data, making the inference of such pMHC-specific models challenging. To address this, it would be appealing to use generalizable rules as priors when inferring the selection landscape for a given pMHC.

One approach involves using the pMHC-specific selection factors $Q_\pi(\sigma)$, such that $P_\pi(\sigma) = Q_\pi(\sigma)P(\sigma)$, where $P(\sigma)$ represents the highly nonuniform [31] but largely universal [32] baseline distribution resulting from VDJ recombination [26,33,34]. A naive strategy might average selection factors across pMHCs, $\langle Q_\pi(\sigma) \rangle_{P(\pi)}$ as an initial guess, where $\langle f(x) \rangle_{P(x)}$ denotes the average of $f(x)$ over the distribution $P(x)$. Yet, such a “mean-field” approach would obscure much of the information about TCR-pMHC specificity.

The binding energy landscapes associated with specific ligands peak in different regions of receptor sequence space [Fig. 1(a)], and averaging over these disordered landscapes is expected to result in a flat and noninformative marginal distribution [Fig. 1(a), inset]. Indeed, embeddings from masked language models trained to reproduce overall repertoire statistics demonstrate limited transfer learning capability for predicting TCR-pMHC specificity [20]. Furthermore, selection factors on observed naive or memory TCR repertoires relative to null expectations are largely driven by

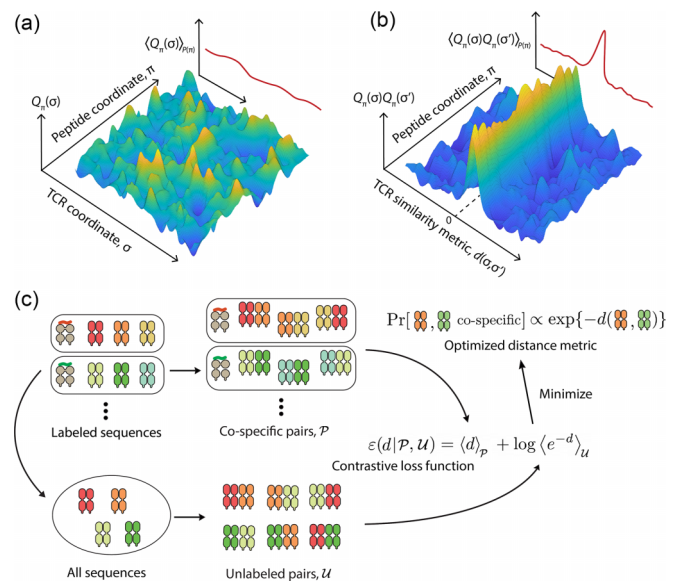


FIG. 1. Contrastive learning of rules that generalize across complex receptor-ligand maps. (a) Cartoon disordered “landscape” of selection factors $Q_\pi(\sigma)$, which describe the varying binding affinities of TCRs σ to ligands π . Averaging selection factors over ligands, $\langle Q_\pi(\sigma) \rangle_{P(\pi)}$, will lead to a largely flat marginal distribution (inset), unless we are able to restrict the average to ligands with highly similar TCR landscapes. (b) Cartoon “landscape” of coselection factors $Q_\pi(\sigma)Q_\pi(\sigma')$, which relate to the probability with which two TCRs are specific to the same ligand, as a function of TCR similarity $d(\sigma, \sigma')$. In contrast to the previous scenario, coselection factors have a nontrivial marginal distribution when averaged across unrelated ligands, $\langle Q_\pi(\sigma)Q_\pi(\sigma') \rangle_{P(\pi)}$ (inset). (c) Sketch of the supervised contrastive learning paradigm. Experimentally determined pairs of TCR sequences binding the same ligand make up the set of cospecific pairs \mathcal{P} , and all pairs of TCR sequences irrespective of specificity make up the set of unlabeled pairs \mathcal{U} . We optimize parameters θ of a family of distance metrics d_θ that minimize a contrastive loss [Eq. (7)]. The loss function minimizes the average distance between pairs in \mathcal{P} while pushing apart unlabeled pairs in \mathcal{U} . Optimized parameters θ^* relate the sequence similarity between a pair of TCRs to the probability with which they share ligand specificity and provide biophysical insights into feature importance.

pMHC-independent constraints on receptor function, such as those required for proper folding [33].

To overcome this problem, we have followed a different approach towards identifying reproducible statistical order across receptor-ligand maps, which rests on analyzing two-point statistics of selection $Q_\pi(\sigma)Q_\pi(\sigma')$ [20,26–28]. The two-point statistics describe the likelihood that receptors σ and σ' are both specific to the same peptide π [Fig. 1(b)]. These coselection factors can have nontrivial averages across unrelated ligands [Fig. 1(b), inset], even where one-point statistics do not. Intuitively, peak positions vary across the selection landscapes, but the size and shape of the peaks might not. Two-point statistics have also been successfully applied in the statistical physics of spin glasses to reveal otherwise hidden statistical order [35].

Using this approach, we previously demonstrated that coselection factors decay predictably with the Levenshtein

distance between TCR sequences, independent of the specific ligand [26]. These findings suggest that generalizable statistical order exists across receptor-ligand maps, which can inform metrics for TCR sequence analysis.

III. SUPERVISED CONTRASTIVE LEARNING AS PSEUDOLIKELIHOOD MAXIMIZATION

A. Derivation of a loss function

We propose fitting available cospecificity data by optimizing parameters θ of a family of TCR sequence similarity metrics $d_\theta(\sigma, \sigma')$. The fitting is based on a supervised contrastive loss function, similar to others used in the machine learning literature [36]. In the following, we provide a derivation of this loss function that clarifies its connection to pseudolikelihood maximization—a technique used by statistical physicists to infer Potts models for protein families [37].

The experimental data take the form of a collection of TCRs $\{\sigma_1, \sigma_2, \dots, \sigma_N\}$, where each sequence σ_i is associated with a cognate pMHC π_i [Fig. 1(c)]. To find an optimal metric, we propose maximizing the pseudolikelihood of all observed pairs of cospecific TCRs:

$$\mathcal{L} = \prod_{\substack{i \neq j \\ \pi_i = \pi_j}} Q_{\pi_i}(\sigma_i) Q_{\pi_i}(\sigma_j) P(\sigma_i) P(\sigma_j). \quad (1)$$

We have shown previously that coselection factors between randomly chosen pairs of sequences at a given Levenshtein distance $d_{LD}(\sigma, \sigma') = \Delta$,

$$\langle Q_\pi(\sigma) Q_\pi(\sigma') \rangle_{P(\sigma, \sigma' | d_{LD}(\sigma, \sigma') = \Delta)}, \quad (2)$$

decay exponentially at small distances Δ with a typical, ligand-independent lengthscale [26]. In fitting a metric to maximize the pseudolikelihood defined in Eq. (1), we thus focus on capturing the average behavior of coselection factors across pMHCs,

$$Q_{\pi_i}(\sigma) Q_{\pi_i}(\sigma') \approx Q(\sigma, \sigma') \equiv \langle Q_\pi(\sigma) Q_\pi(\sigma') \rangle_{P(\pi)}, \quad (3)$$

according to an exponential ansatz

$$Q(\sigma, \sigma') \propto e^{-d_\theta(\sigma, \sigma')}. \quad (4)$$

Although outside the scope of our study, it is clear that this metric could be improved for prediction on specific targets by exploiting ligand-dependent variation in these rules [27,28], as we have previously demonstrated for protein language models [20]. Here, inference is instead restricted by design to cospecificity rules that generalize across ligands.

To complete the specification of the optimization problem, it remains to determine the proportionality constant in Eq. (4). We use a normalization condition $\langle Q(\sigma, \sigma') \rangle_{P(\sigma)P(\sigma')} = 1$, such that

$$Q(\sigma, \sigma') = \frac{e^{-d_\theta(\sigma, \sigma')}}{\langle e^{-d_\theta(\sigma, \sigma')} \rangle_{P(\sigma)P(\sigma')}}. \quad (5)$$

With sufficient data, the average in the denominator can be approximated by summation over all pairs of TCRs in the training data. We call this set of pairs the unlabeled set, $\mathcal{U} = \{(\sigma_i, \sigma_j) | i \neq j, \forall i, j\}$, as it ignores the associated pMHC labels. This means approximating the partition function with respect to $P(\sigma)P(\sigma')$ by an average with respect

to the empirical distribution of unlabeled pairs $\mathcal{U}(\sigma, \sigma') = |\mathcal{U}|^{-1} \sum_{(s, s') \in \mathcal{U}} \delta_{\sigma, s} \delta_{\sigma', s'}$. Here, $|\mathcal{U}|$ denotes the cardinality of the unlabeled set, and $\delta_{x, y}$ denotes the Kronecker delta, $\delta_{x, y} = 1$ for $x = y$ and $\delta_{x, y} = 0$ otherwise. We note that an alternative approach would be to calculate the average only with respect to pairs of TCRs with differing pMHC specificity or by using a large set of independently acquired TCRs, say from an unsorted blood sample. However, in line with prior work [20,38], we found that the inclusion of cospecific pairs in the evaluation of the denominator has a regularizing effect on the inference procedure.

To simplify notations, we similarly define the set of positive pairs, $\mathcal{P} = \{(\sigma_i, \sigma_j) | \pi_i = \pi_j, \forall i, j\}$, and the empirical probability of observing TCR pairs (σ, σ') in the set of cospecific pairs, $\mathcal{P}(\sigma, \sigma') = |\mathcal{P}|^{-1} \sum_{(s, s') \in \mathcal{P}} \delta_{\sigma, s} \delta_{\sigma', s'}$. Plugging all definitions into Eq. (1) allows us to write the negative log-pseudolikelihood per positive pair as

$$-\frac{\log \mathcal{L}}{|\mathcal{P}|} = \langle d_\theta(\sigma, \sigma') \rangle_{\mathcal{P}(\sigma, \sigma')} + \log \langle e^{-d_\theta(\sigma, \sigma')} \rangle_{\mathcal{U}(\sigma, \sigma')} + c, \quad (6)$$

where $c = -(\log P(\sigma)P(\sigma'))_{\mathcal{P}(\sigma, \sigma')}$ is a term that is independent of the distance metric d_θ . Optimal parameters θ^* maximize the pseudolikelihood [Eq. (1)], or equivalently, minimize terms in the negative log-likelihood dependent on d_θ , which yields the following loss function:

$$\varepsilon(\theta | \mathcal{P}, \mathcal{U}) = \langle d_\theta(\sigma, \sigma') \rangle_{\mathcal{P}(\sigma, \sigma')} + \log \langle e^{-d_\theta(\sigma, \sigma')} \rangle_{\mathcal{U}(\sigma, \sigma')}. \quad (7)$$

Figure 1(c) summarizes in a flowchart how the sets \mathcal{P} and \mathcal{U} of cospecific and unlabeled sequence pairs are constructed.

B. Analytical solution for a simple case

To build intuition, we minimize Eq. (7) in closed form for a simple distance metric. We consider distance metrics defined in terms of a discrete-valued function g that maps each pair of TCRs to a finite set of K integers [e.g., for $K = 2$ a simple case would be $g(\sigma, \sigma') = 1$ if there is an insertion or deletion and $g(\sigma, \sigma') = 0$ for an amino-acid substitution]. Given the integer-valued function g , a family of distance metrics with trainable parameters can be defined by

$$d_g(\sigma, \sigma') = \sum_{i=1}^K \theta_i \delta_{g(\sigma, \sigma'), i}, \quad (8)$$

where $\theta_1, \theta_2, \dots, \theta_K$ are parameters assigned to the K possible value of g . Substituting Eq. (8) into the loss function [Eq. (7)] yields

$$\varepsilon(\theta | \mathcal{P}, \mathcal{U}) = \sum_{i=1}^K \theta_i \langle \delta_{g(\sigma, \sigma'), i} \rangle_{\mathcal{P}(\sigma, \sigma')} + \log Z, \quad (9)$$

where $Z = \langle e^{-\sum_i \theta_i \delta_{g(\sigma, \sigma'), i}} \rangle_{\mathcal{U}(\sigma, \sigma')}$ is a normalizing constant. At the optimal θ_i^* , we have $\partial \varepsilon / \partial \theta_i = 0$. Taking the partial derivative of Eq. (9) with respect to θ_i yields

$$\frac{\partial \varepsilon}{\partial \theta_i} = \langle \delta_{g(\sigma, \sigma'), i} \rangle_{\mathcal{P}(\sigma, \sigma')} - \frac{e^{-\theta_i}}{Z} \langle \delta_{g(\sigma, \sigma'), i} \rangle_{\mathcal{U}(\sigma, \sigma')}. \quad (10)$$

Solving for θ_i^* by setting Eq. (10) to zero, we derive

$$\theta_i^* = -\log \frac{\langle \delta_{g(\sigma, \sigma'), i} \rangle_{\mathcal{P}(\sigma, \sigma')}}{\langle \delta_{g(\sigma, \sigma'), i} \rangle_{\mathcal{U}(\sigma, \sigma')}} - \log Z. \quad (11)$$

This result shows that minimizing the contrastive loss function [Eq. (7)] sets parameters based on the log-ratio between observed occurrences of particular features among positive pairs and a null expectation calculated across unlabeled pairs.

This analytical solution shows that contrastive learning using Eq. (7) is closely connected to how BLOSUM and related evolutionary amino-acid substitution matrices are traditionally defined in the bioinformatics literature [24,39]. These matrices are also constructed by calculating log-odds ratios between observed and expected frequencies of amino acid pairs. However, contrastive learning extends this by accommodating more complex distance metrics that incorporate multiple substitution properties. Furthermore, empirical frequencies across unlabeled TCR pairs are used in the denominator, rather than assuming factorization as in BLOSUM. This is a more appropriate choice for TCRs, where VDJ recombination introduces correlations between sites even in the naive repertoire [31].

IV. DATA PREPARATION

To learn rules that generalize to unseen pMHCs, we curated the training data to eliminate confounding factors and to reduce the complexity of the problem. We decided to restrict training to a single dataset of TCRs with experimentally annotated pMHC specificity to exclude confounding by data origin. Consequently, we used a large-scale dataset on TCRs with specificity to SARS-CoV-2 peptides obtained using a multiplexed assay called MIRA [11]. From these data we only used pairs of TCR CDR3 β sequences that differ by a single edit step to maximize the signal-to-noise ratio given potentially misannotated sequences [26]. Moreover, in both training and testing, we only considered pairs of sequences with identical V genes to ensure that learned rules only reflected amino-acid substitutions in the CDR3 β chain. We also excluded the (very small) number of pairs involving substitutions of cysteine to reduce overfitting, due to the rarity of cysteines within the CDR3 junction following negative thymic selection [33].

To reduce overfitting to less commonly observed sequence lengths, we restricted our analysis to TCR CDR3 β sequences with lengths between 11 and 15. These represented 72% of the sequences in the MIRA database and 76% of sequences in the VDJdb database. From the MIRA database, we identified 26 sets of TCRs specific to unique pMHCs, each containing at least 1000 TCR sequences. To mitigate the impact of varying sizes of specificity groups, we equalized all group sizes by subsampling without replacements to 1000 sequences per pMHC. To still learn on all available data, we repeated the procedure 15 times and averaged performance across these training batches. Among the 26 000 TCRs included per training batch, we identified on average ≈ 6000 pairs of single-edit TCRs.

We evaluated the inferred metric using TCRs specific to hold-out pMHCs from the MIRA database. To this end, we collated sets of cospecific sequence pairs from 24 specificity groups, each containing between 400 and 999 TCR sequences, that were excluded during training. We again used subsampling to address class imbalance, generating 15 sets of 400 sequences for each specificity group. Furthermore, to assess generalization to TCR pairs differing by multiple edits, we

identified TCR pairs with edit distances ranging from 1 to 6 in the same dataset.

Using these data, we chose to test the ability of metrics to discriminate between cospecific (i.e., $\in \mathcal{P}$) and cross-specific TCR pairs [i.e., $\in \mathcal{C} = \{(\sigma_i, \sigma_j) | \pi_i \neq \pi_j, \forall i, j\}$]. We note that this task represents an instance of positive-unlabeled classification [40], as cross-specific pairs are not tested experimentally for cospecificity. We have shown previously that about one in ten single edits maintains cospecificity [26], setting an upper bound on achievable classification accuracy. Nonetheless, we chose this task for its simplicity. Additionally, when considering TCR pairs differing by multiple edits, the prevalence of cospecific pairs decreases, reducing the impact of this limitation.

To externally validate the predictive power of the inferred rules, we used data from VDJdb [41], excluding any pMHCs present in the MIRA dataset. In addition to TCR β sequences, VDJdb includes TCR α sequences for certain pMHCs, which we used to further test the biological generalizability of the inferred cospecificity rules. For the β chain validation task, we selected TCRs from 131 specificity groups, each containing at least 15 unique sequences. As before, this threshold was chosen to maximize the number of cospecific pairs included in the analysis after balancing the data across groups. For the α chain validation task, we identified 49 specificity groups meeting the same threshold of 15 sequences. Consistent with our previous methodology, we employed subsampling and identified TCR pairs with edit distances ranging from 1 to 6 across 15 balanced testing batches.

V. INFERENCE OF OPTIMAL SEQUENCE SIMILARITY METRICS

A. Overview, notations, and baseline metric

To provide proof-of-concept for the applicability of our framework to this dataset, we applied our inference procedure to optimize parameters within a family of simple alignment-based distance metrics. Alignment-based metrics of TCR sequence similarity have been widely used to predict TCR cospecificity [3,4,42,43]. Among these, one of the most widely used metrics is TCRdist [3], which calculates distances between TCRs based on BLOSUM62-derived amino-acid dissimilarity penalties. TCRdist furthermore optionally trims CDR3 sequences to exclude residues at the edges from the mismatch calculation, as a simple form of nonuniform weighting of mismatches according to position inspired by structural evidence showing that these flanking residues rarely make direct contact with the peptide [3,43]. Building on the principles of TCRdist, we focused on the inference of parameters relating to two critical physical factors influencing the effects of changes in the TCR hypervariable region: the identity of the substituted amino acids and the position of the substitution.

Specifically, we define distance metrics in terms of edit steps, which characterize the difference between a pair of TCR sequences σ and σ' by a series of substitutions, deletions, and insertions [Fig. 2(a)]. The simplest among these metrics, the total unweighted number of edit steps required to transform σ to σ' , is known as the Levenshtein distance [44]. In the following, we describe the i th edit step between σ and σ'

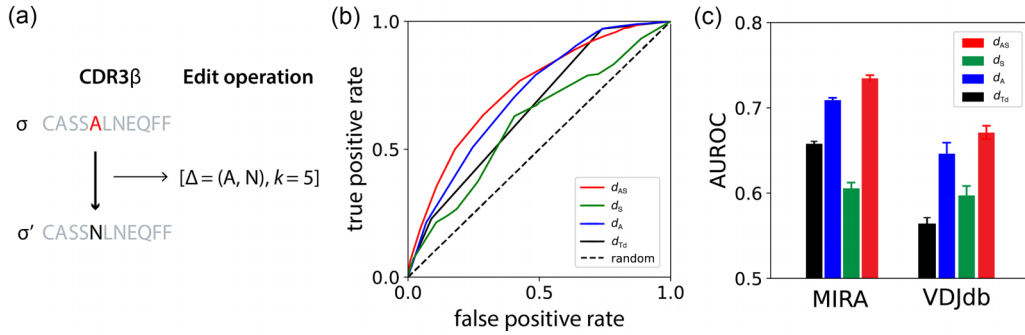


FIG. 2. Learning of a CDR3 β distance metric that generalizes to unseen ligands. (a) Distance metrics are defined in terms of the edit steps between two CDR3 β sequences. Each edit step includes the length of sequence ℓ , the position of substitution κ , and the identity of substitution Δ . (b) Receiver Operating Characteristics (ROC) curves for identifying cospecific relative to cross-specific [i.e., $\in \mathcal{C} = \{(\sigma_i, \sigma_j) | \pi_i \neq \pi_j, \forall i, j\}$] TCR CDR3 β pairs from the MIRA test set that differ by a single edit using different TCR similarity metrics. The metrics include TCRdist (d_{Td} , black) as a baseline, as well as a series of metrics optimized using the mathematical framework proposed in this study. These metrics involve substitution type (which amino acids are substituted) and/or substitution site (which position along the CDR3 sequence is substituted). Specifically, the curves show metrics using only type (d_A , blue), only site (d_S , green), or using both jointly (d_{AS} , red). Curves are averaged over 15 subsampled test sets, each equally balanced across pMHC groups. (c) Areas under the ROC (AUROC) for the MIRA test set and an external validation set (VDJdb). Error bars indicate the standard error over resampled testing batches. Note that due to cross-reactivity some cross-specific pairs also exhibit cospecificity, which bounds achievable AUROCs below 1.

by $\Delta_i(\sigma, \sigma')$, the identity of substituted amino acids, and by $k_i(\sigma, \sigma')$, the substitution position along the TCR sequence [Fig. 2(a)].

In this notation, the baseline metric, TCRdist without trimming, is given by $d_{Td}(\sigma, \sigma') = \sum_i \tilde{d}_{Td}(\Delta_i)$, where $\tilde{d}_{Td}(\Delta_i) = \min(4, 4 - B_{62}(\Delta_i))$ if the i th edit step is an amino-acid substitution, and $\tilde{d}_{Td}(\Delta_i) = 8$ if the i th edit step is an insertion/deletion (see Fig. S1A of the Supplemental Material [45]). Here, B_{62} denotes the BLOSUM62 matrix. In the remainder of this section, we use d_{Td} as a baseline for comparisons.

B. Learning an amino-acid substitution matrix

To test whether an amino-acid substitution matrix directly learned from TCR data would outperform the BLOSUM-derived TCRdist score, we used our framework to train the following simple metric:

$$d_A(\sigma, \sigma') = \sum_i A(\Delta_i), \quad (12)$$

where A is a matrix of amino-acid substitution weights (Fig. S1B). We included a blank (“—”) in the amino-acid alphabet to treat insertion/deletion as a substitution involving “—”.

To evaluate the performance of d_A , we generated pairs of TCR CDR3 β sequences from the MIRA database that differed by a single edit step and belonged to a specificity group not included in the training data. The performance of distance metrics was quantified using a cospecific pair identification test, in which cospecific TCRs were treated as positive pairs, while TCR pairs across different specificity groups were treated as negative pairs as described before. The receiver operating characteristic (ROC) curves for different models show that d_A outperforms d_{Td} [Fig. 2(b)]. Furthermore, the learned substitution matrix also better classifies cospecific TCR CDR3 β sequence pairs from the VDJdb database [Fig. 2(c)], as assessed by the area under the receiver

operating curve (AUROC). Both tasks involve TCRs specific to unseen pMHCs, which in the case of VDJdb additionally originate from pathogens other than SARS-CoV-2, which was used during training.

C. Substitution position influences TCR cospecificity

We next investigated whether the position of substitution (independently of the identity of the substituted amino acids) determines the likelihood of TCR cospecificity, as suggested by structures [4,25]. To address this question, we trained a distance metric that considers only the position of substitution:

$$d_S(\sigma, \sigma') = \sum_i S(k_i, \ell), \quad (13)$$

where S represents a matrix of site weights dependent on k_i , the position of substitution, and $\ell(\sigma, \sigma') = \max(|\sigma|, |\sigma'|)$, the aligned length of the TCR pair (Fig. S2A).

Results on the testing task show that d_S has some predictive power on unseen pMHCs [Fig. 2(b)], and even outperforms d_{Td} on the validation task [Fig. 2(c)]. These findings confirm that location matters for the impact of a substitution on TCR specificity.

D. Joint inference of site- and identity-dependent substitution weights

Having established that both the identity of substituted amino acids and the position of substitution influence the probability of cospecificity in a pair of TCRs, we next inferred a model accounting for both factors jointly. Incorporating multiple factors into phenomenological distance metrics is challenging, as it requires ad hoc choices for setting the relative weight given to each factor. In contrast, in our framework, different parameters are jointly optimized during the training process to disentangle their relative influence.

To jointly optimize site and identity-dependent weights while keeping the number of parameters manageable, we




FIG. 3. Comparison of site-dependent weights with TCR-pMHC contact probabilities reveals the importance of noncontact sites. (a) Site weights, P_S , learned from joint optimization of substitution type-dependent and site-dependent weights (d_{AS}). Sites that lack sufficient substitution statistics are shown in gray. (b) Average contact probability in TCR-pMHC crystal structures (see Methods). Sites along the CDR3 β sequence that never contact the peptide in known structures are shown in gray. (c) Scatter plot showing the correlation of both quantities across sites.

defined a factorized family of distance metrics,

$$d_{AS}(\sigma, \sigma') = - \sum_i \log[1 - (1 - e^{-S(k_i, \ell)})(1 - e^{-A(\Delta_i)})], \quad (14)$$

where S and A are, respectively, the co-optimized substitution weights that depend on the position of substitution and the identity of the substitution. At first order, this expression simplifies to the product of weights, $d_{AS}(\sigma, \sigma') \sim \sum_i S(k_i, \ell)A(\Delta_i)$, but inspired by prior models of receptor-ligand interactions [17,26,46], one can construct a simple toy model to give Eq. (14) a more biophysical interpretation. We define $P_S(k, \ell)$ as the probability that a residue at site k on a TCR of length ℓ contributes to pMHC specificity, and $P_A(\Delta)$ as the probability that cospecificity is maintained at a given site following a substitution Δ . Most naively, $P_S(k, \ell)$ can be interpreted as a contact probability between TCR residue and target peptide, but the importance of a site might also reflect MHC contact or allostery. Assuming further that substitution penalties are only incurred at sites important to binding and that P_S and P_A are statistically independent, we can express coselection factors as

$$Q_{AS}(\sigma, \sigma') \propto \prod_i \{1 - P_S(k_i, \ell)[1 - P_A(\Delta_i)]\}. \quad (15)$$

We set $Q_{AS}(\sigma, \sigma') \propto e^{-d_{AS}(\sigma, \sigma')}$ and parametrize the probabilities using substitution weights similar to those used in the previous single-factor optimization: $P_S(k_i, \ell) = 1 - e^{-S(k_i, \ell)}$ and $P_A(\Delta_i) = e^{-A(\Delta_i)}$. Using these definitions, it follows that Eq. (15) is equivalent to the family of distance metrics defined in Eq. (14).

We found that by simultaneously accounting for both the position and identity of substitutions, d_{AS} outperforms all other distance metrics [Fig. 2(b)]. This improved performance extends to the external validation on data from VDJdb [Fig. 2(c)]. Further stratification of the testing data reinforces that the learned rules generalize broadly. They remain predictive even when tested on peptides differing by ≥ 6 amino

acids from the nearest training example, as well as when filtering TCRs by sequence similarity and when stratifying by HLA (Fig. S5). Generalization to such highly dissimilar peptides has remained out of reach of previous sequence-based machine learning approaches [18,21], providing empirical support for the theoretical intuition developed in Sec. II regarding the greater universality of cospecificity rules.

VI. BIOPHYSICAL INTERPRETATION OF INFERRED PARAMETERS

What physical properties govern TCR specificity? Having established that the position and identity of substitutions determine the likelihood of TCR cospecificity, we turned our attention to the biophysical interpretation of the optimized parameters of the jointly optimized d_{AS} [Eq. (14)].

To judge how reliably parameters could be inferred from the available training data, we calculated their coefficients of variation across resampled training sets. We found them to be modest (Figs. S3A,B and Figs. S4A,B), particularly for more commonly observed substitutions (Figs. S3C and S4C). We expect this variability to further decrease as more data become available. This would potentially boost the correlation coefficients between fitted site and substitution weights and biophysical parameters, but some broad insights already emerge from the analysis of the available data.

To test whether site weights P_S , where higher values signify increased impact on cospecificity, are in line with their intuitive role as a proxy for contact probabilities, we calculated contact probabilities from TCR-pMHC complexes deposited in the Structural T Cell Receptor Database [47] [Fig. 3(b); see Methods for details]. We found that P_S is moderately correlated with these contact probabilities, with a Spearman's rank correlation of $r = 0.47$ [Fig. 3(c)]. Weights S learned by single-factor optimization (Fig. S2A) were strongly correlated with jointly optimized weights S (Fig. S2C), but showed a marginally weaker correlation of $r = 0.45$ with contact probabilities (Fig. S2B). In contrast to contact probabilities, the

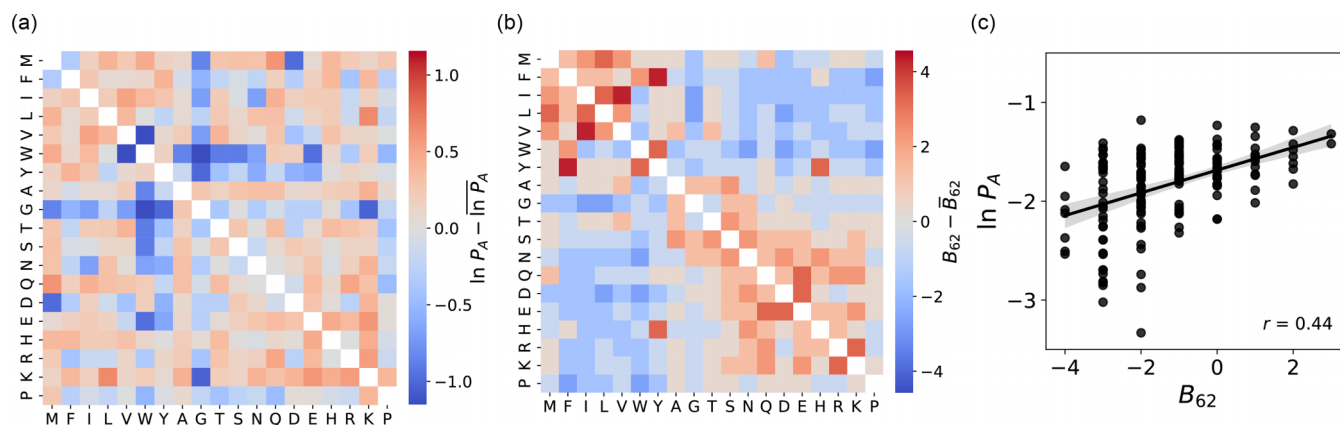


FIG. 4. Optimized amino-acid substitution matrix differs from BLOSUM. (a) Probability of maintaining cospecificity, P_A , learned from joint optimization of substitution type-dependent and site-dependent weights (d_{AS}). Note that cysteine was removed from the optimization due to a lack of substitution statistics in CDR3 β sequences. (b) BLOSUM62 amino-acid substitution matrix. (c) Comparison of learned P_A and the BLOSUM62 matrix.

optimal metric still heavily weights the edges of the CDR3 loop. This implies that substitutions of residues that do not make direct contact with the peptide substantially impact TCR specificity and should be taken into account for accurate cospecificity prediction.

We next compared the learned amino-acid substitution matrix $P_A(\Delta)$ with prior expectations. Our results show that as expected, TCR pairs differing by insertions or deletions are much less likely to be cospecific compared to those differing by amino-acid substitutions (Fig. S1C). We further compared $\log P_A$ [Fig. 4(a)] and the widely used BLOSUM62 matrix B_{62} [Fig. 4(b)]. Both matrices are moderately correlated with a Spearman's rank correlation coefficient of

$r = 0.44$ [Fig. 4(c)]. The jointly optimized amino-acid substitution matrix \mathcal{A} was also correlated with the truncated BLOSUM62 variant used by TCRdist (Fig. S1E), and the matrix A obtained by single-factor optimization (Fig. S1F). Single-factor optimization led to marginally weaker correlations with BLOSUM-based scores (Fig. S1A).

What factors drive observed differences between \mathcal{A} and B_{62} ? To answer this question, we identified the physical properties of amino acids that predict most closely the substitution matrices. For a given property $h(n)$ of amino acid n , we quantified the contribution of h to the substitution matrices by calculating the matrix of absolute differences in h , defined as $\mathcal{H}(n, m) = |h(n) - h(m)|$. We then calculated

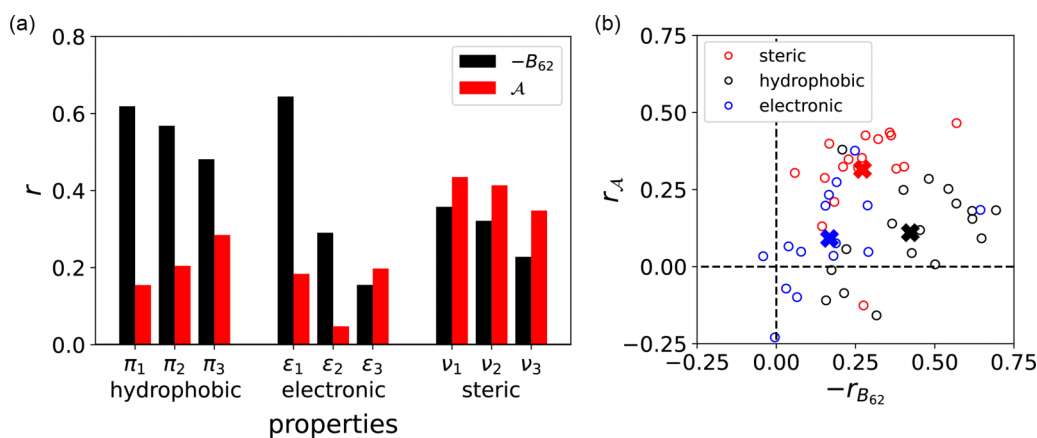


FIG. 5. Physical correlates of the optimized amino-acid substitution matrix highlight the contribution of shape to TCR specificity. (a) Spearman's rank correlation coefficient (r) of substitution matrices with matrices constructed from pairwise absolute difference in physical properties of amino acids. Selected hydrophobic, electronic, and steric properties were compared. Hydrophobic properties: partition coefficient in octanol/water (π_1), solvation free energy (π_2), and dG of transfer from organic solvent to water (π_3). Electronic properties: polarity (ϵ_1), pK of COOH on C_α (ϵ_2), and pK of NH_2 on C_α (ϵ_3). Steric properties: average volume of buried residue (v_1), normalized van der Waals volume (v_2), and STERIMOL maximum width of the side chain (v_3). (b) Comparison of the correlation coefficients of both substitution matrices with those obtained from a panel of 50 amino-acid properties (for the complete list, see the Supplemental Material). Centroids for each property class are shown as "x."




FIG. 6. Learned metrics generalize to pairs of TCRs with a range of sequence similarities. AUROC of baseline method TCRdist (d_{TD} , black) and inferred site-specific metric (d_S , green), amino-acid substitution metric (d_A , blue), and jointly learned metric (d_{AS} , red) for edit distance 1–6 TCR pairs from different datasets. (a) MIRA test set. (b) VDJDdb CDR3 β validation set. (c) VDJDdb CDR3 α validation set. Error bars show the standard error over resampled test batches.

Spearman’s rank correlation coefficient between the entries of \mathcal{H} and the entries of \mathcal{A} and B_{62} for nine amino-acid properties.

As shown in Fig. 5(a), we find that B_{62} is strongly dependent on a number of hydrophobic and electronic properties of amino acids. In contrast, these properties are less correlated with the learned substitution matrix \mathcal{A} . Instead, \mathcal{A} correlates most with the steric similarity of substituted amino acids. For instance, many of the most strongly penalized substitutions involve replacing glycine, the amino acid with the smallest and most conformationally flexible side chain. These results are compatible with our prior findings using amino-acid properties to determine optimal reduced alphabets [28]. The importance of steric properties in determining the specificity of TCRs was further confirmed by analyzing a larger panel of 50 amino-acid properties [Fig. 5(b)] taken from Ref. [48] (listed in Table S1).

To test whether the observed decrease in the significance of hydrophobic properties—and the corresponding increase in the significance of steric properties—extends to other classes of protein-protein interactions, we analyzed thermodynamic measurements from the SKEMPI database [49]. Specifically, we examined which physical properties of residues (Table S1) most strongly correlate with the change in binding free energy $\Delta\Delta G$ upon mutation. As shown in Fig. S6, steric properties are most predictive for both TCR-pMHC and antibody-antigen interactions. In contrast, for protease-inhibitor interactions, hydrophobicity emerged as the most significant factor. These results are compatible with the hypothesis that steric effects may be particularly important in interactions involving unstructured or flexible regions—such as the CDR loops in TCRs and antibodies—whereas hydrophobicity may dominate in interactions mediated by well-structured binding pockets, such as those found in proteases.

VII. EXTRAPOLATION OF LEARNED RULES

Having learned biophysically interpretable local rules based on pairs of single-edit TCR β sequences, we asked whether these rules can be extrapolated to more dissimilar TCR pairs and the TCR α chain. To quantify this, we tested

the ability of trained metrics to identify cospecific sequence pairs among TCRs with an increasing number of edits on the β chain [Figs. 6(a) and 6(b)] and the α chain [Fig. 6(c)]. This approach allowed us to evaluate the robustness and applicability of the metric across a broader range of sequence variations, providing insights into whether the local substitution rules remain effective when sequence dissimilarity increases.

We find that d_{AS} and d_S consistently outperform d_{TD} across all distances [Figs. 6(a) and 6(b)]. These metrics, which were optimized to take into account amino-acid substitution type (d_A and d_{AS}), retain strong predictive power for TCR pairs with edit distances up to three to four edit steps. When applied to VDJDdb, classification performance for certain metrics even increases for sequences that differ by two or three edits [Fig. 6(b)], as might be expected within the range of applicability of an independent substitution model. This relatively high predictability at moderate edit distances may, in part, reflect compound effects of multiple substitutions that make it easier to distinguish true positives from negatives. Additionally, it is important to interpret this performance in light of the labeling strategy: because we label TCR pairs from different specificity groups as negative examples, there is a greater chance of mislabeling for more similar TCRs—i.e., pairs that differ by only one or two substitutions may actually be cross-reactive, artificially decreasing the performance. At larger edit distances, mislabeling of negative examples is no longer significant, as cospecific sequences become very rare. The decline in performance at even larger edit distances is likely due to the growing influence of higher-order features—such as multiple structurally distinct binding solutions [26]—that are not captured by a simple independent substitution model. These findings demonstrate the strength of basic distance metrics for refining notions of similarity for TCRs, while also highlighting their limitations when comparing receptors that are highly dissimilar.

We furthermore tested the generalization of the metrics to the α chain using data from VDJDdb. We find that the learned amino-acid substitution matrix improves upon TCRdist scoring across a range of edit distances [Fig. 6(c)]. These results indicate that some of the local cospecificity rules are common to both α and β chains.

VIII. DISCUSSION

Despite rapid growth in TCR-pMHC data, only a tiny fraction of the vast diversity of pairings has been measured. This sparsity complicates the discovery of generalizable rules. In this study, we addressed this problem by developing a pseudolikelihood maximization framework for optimizing notions of TCR sequence similarity. We demonstrate that it is possible to use this framework to learn rules that robustly generalize to pMHCs not used during training. Our findings highlight that, while TCR-pMHC interactions are globally complex, local sequence-based rules effectively predict cospecificity, offering a practical advance for deorphanizing TCRs of unknown specificity.

Our results challenge the use of the BLOSUM62 amino-acid substitution matrix for TCR sequence analysis. These insights confirm the independent findings of Postovskaya *et al.* [50] using alternative approaches and data, further underscoring the need for tailored distance matrices. Specifically, we observed that hydrophobicity was a less significant factor than previously assumed, while steric similarity emerged as a more meaningful—albeit moderate—predictor of substitution impact. This shift in emphasis aligns with the classical immunological “shape space” concept [51], which posits that shape complementarity is central to receptor-ligand interactions. Interestingly, while hydrophobicity showed limited predictive value for two-point statistics in our study, hydrophobic CDR3 residues are important in promoting the development of self-reactive T cells [52], highlighting a context-dependent role. Comparison of position-specific terms in more complex models with structural data further highlighted the importance of non-peptide-contact sites within the CDR3 to binding specificity. Such a dependence on context might limit the predictive performance of TCR metrics using trimmed sequences or machine learning approaches solely taking into account contacting residues [53], and it is compatible with the presumptive role of all residues in determining overall loop geometry.

The application of our optimization framework was limited by the available sequence data. To make robust inferences possible, we decided to simplify the problem setting and restrict the family of metrics used for inference. We trained on CDR3 β sequences of restricted lengths and only considered sequence pairs with matching V genes. Moreover, we only considered simple metrics with a limited number of free parameters. We assumed, for example, a factorization of site-dependent and substitution-dependent terms. With increasing data availability, including more diverse data on paired-chain TCR $\alpha\beta$ sequences and their pMHC ligands, our framework can be applied to learn more sophisticated distance metrics.

Ultimately, we envisage that it will be possible to derive a machine-learning optimized mlTCRdist metric, which can act as a drop-in replacement for TCRdist. Additionally, we hope that the biophysical insights, robust learning framework, and data curation strategy introduced in this study can also be used

to overcome current bottlenecks in the supervised contrastive training of TCR representations using deep learning architectures [20,54]. Better TCR representations could significantly enhance TCR clustering for biomarker and metaclonotype discovery [43]. Improved understanding of the biophysical determinants of TCR specificity could also advance rational TCR design for therapeutic applications [55].

IX. METHODS

A. Training procedure

Parameters were optimized by gradient descent, using the Adam algorithm as implemented in PyTorch [56]. To reduce sampling bias, the summed loss over 15 batches with equal representation of different pMHCs was simultaneously minimized, i.e., $\varepsilon(d) = \sum_{i=1}^{15} \varepsilon(d|\mathcal{P}_i, \mathcal{U}_i)$. All distance metrics were trained with a learning rate of 10^{-3} , where 3000 steps were taken to train d_S , 5000 steps for d_A , and 12 000 steps for d_{AS} .

B. Calculation of contact probabilities

To obtain TCR-peptide contact information, we utilized 179 publicly available TCR-pMHC complexes from the Structural T Cell Receptor Database [47]. The pairwise distances between each residue in the TCR CDR3 β loop and each peptide residue were calculated as the minimum distance between nonhydrogen atoms, as previously described [25]. Residues were defined as being in contact if their pairwise distance was less than or equal to 5 Å.

ACKNOWLEDGMENTS

A.G.T.P. was supported by Natural Sciences and Engineering Research Council of Canada (NSERC). Y.N. and M.M. were supported by Cancer Research UK studentships under Grants No. BCCG1C8R and No. A29287, respectively. J.H. was supported in part by an award (to A.T.-M.) of the National Institute for Health Research Biomedical Research Centre at NIHR UCLH BRC. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR, or the Department of Health. N.S.W. was supported in part by the National Science Foundation, through the Center for the Physics of Biological Function (PHY-1734030). A.T.-M. was supported in part by funding from the Royal Free Charity. This research was supported in part by the National Science Foundation Grant No. PHY-2309135 and the Gordon and Betty Moore Foundation Grant No. 2919.02 to the Kavli Institute for Theoretical Physics (KITP). The authors acknowledge helpful discussions with Gian Marco Visani and Antonio Matas Gil.

DATA AVAILABILITY

The data that support the findings of this article are openly available [57].

[1] M. M. Davis and P. J. Bjorkman, The T cell receptor genes and T cell recognition, *Nature (London)* **334**, 395 (1988).

[2] H. Chi, M. Pepper, and P. G. Thomas, Principles and therapeutic applications of adaptive immunity, *Cell* **187**, 2052 (2024).

- [3] P. Dash, A. J. Fiore-Gartland, T. Hertz, G. C. Wang, S. Sharma, A. Souquette, J. C. Crawford, E. B. Clemens, T. H. O. Nguyen, K. Kedzierska, N. L. La Gruta, P. Bradley, and P. G. Thomas, Quantifiable predictive features define epitope-specific T cell receptor repertoires, *Nature (London)* **547**, 89 (2017).
- [4] J. Glanville, H. Huang, A. Nau, O. Hatton, L. E. Wagar, F. Rubelt, X. Ji, A. Han, S. M. Krams, C. Pettus, N. Haas, C. S. L. Arlehamn, A. Sette, S. D. Boyd, T. J. Scriba, O. M. Martinez, and M. M. Davis, Identifying specificity groups in the T cell receptor repertoire, *Nature (London)* **547**, 94 (2017).
- [5] C. S. Dobson, A. N. Reich, S. Gaglione, B. E. Smith, E. J. Kim, J. Dong, L. Ronsard, V. Okonkwo, D. Lingwood, M. Dougan, S. K. Dougan, and M. E. Birnbaum, Antigen identification and high-throughput interaction mapping by reprogramming viral entry, *Nat. Methods* **19**, 449 (2022).
- [6] A. V. Joglekar, M. T. Leonard, J. D. Jeppson, M. Swift, G. Li, S. Wong, S. Peng, J. M. Zaretsky, J. R. Heath, A. Ribas *et al.*, T cell antigen discovery via signaling and antigen-presenting bifunctional receptors, *Nat. Methods* **16**, 191 (2019).
- [7] S. Sureshchandra, J. Henderson, E. Levendosky, S. Bhattacharyya, J. M. Kastenschmidt, A. M. Sorn, M. T. Mitul, A. Benchorin, K. Batucal, A. Daugherty *et al.*, Tissue determinants of the human T cell receptor repertoire, *bioRxiv* (2024), doi:10.1101/2024.08.17.608295.
- [8] D. V. Bagaev, R. M. A. Vroomans, J. Samir, U. Stervbo, C. Rius, G. Dolton, A. Greenshields-Watson, M. Attaf, E. S. Egorov, I. V. Zvyagin, N. Babel, D. K. Cole, A. J. Godkin, A. K. Sewell, C. Kesmir, D. M. Chudakov, F. Luciani, and M. Shugay, VDJdb in 2019: Database extension, new analysis infrastructure and a T-cell receptor motif compendium, *Nucl. Acids Res.* **48**, D1057 (2020).
- [9] N. Tickotsky, T. Sagiv, J. Prilusky, E. Shifrut, and N. Friedman, McPAS-TCR: A manually curated catalogue of pathology-associated T cell receptor sequences, *Bioinformatics* **33**, 2924 (2017).
- [10] R. Vita, S. Mahajan, J. A. Overton, S. K. Dhanda, S. Martini, J. R. Cantrell, D. K. Wheeler, A. Sette, and B. Peters, The immune epitope database (IEDB): 2018 update, *Nucl. Acids Res.* **47**, D339 (2018).
- [11] S. Nolan, M. Vignali, M. Klinger, J. N. Dines, I. M. Kaplan, E. Svejnova, T. Craft, K. Boland, M. W. Pesesky, R. M. Gittelman *et al.*, A large-scale database of T-cell receptor beta sequences and binding associations from natural and synthetic exposure to SARS-CoV-2, *Front. Immunol.* **16**, 1488851 (2025).
- [12] S. Gielis, P. Moris, W. Bittremieux, N. De Neuter, B. Ogunjimi, K. Laukens, and P. Meysman, Detection of enriched T cell epitope specificity in full T cell receptor sequence repertoires, *Front. Immunol.* **10**, 2820 (2019).
- [13] X. Lin, J. T. George, N. P. Schafer, K. Ng Chau, M. E. Birnbaum, C. Clementi, J. N. Onuchic, and H. Levine, Rapid assessment of T-cell receptor specificity of the immune repertoire, *Nat. Comput. Sci.* **1**, 362 (2021).
- [14] D. S. Fischer, Y. Wu, B. Schubert, and F. J. Theis, Predicting antigen specificity of single T cells based on TCR CDR3 regions, *Mol. Syst. Biol.* **16**, e9416 (2020).
- [15] E. Jokinen, J. Huuhtanen, S. Mustjoki, M. Heinonen, and H. Lähdesmäki, Predicting recognition between T cell receptors and epitopes with TCRGP, *PLoS Comput. Biol.* **17**, e1008814 (2021).
- [16] G. Croce, S. Bobisse, D. L. Moreno, J. Schmidt, P. Guillaume, A. Harari, and D. Gfeller, Deep learning predictions of TCR-epitope interactions reveal epitope-specific chains in dual alpha T cells, *Nat. Commun.* **15**, 3211 (2024).
- [17] A. Wang, X. Lin, K. N. Chau, J. N. Onuchic, H. Levine, and J. T. George, RACER-m leverages structural features for sparse T cell specificity prediction, *Sci. Adv.* **10**, ead10161 (2024).
- [18] B. Meynard-Piganeau, C. Feinauer, M. Weigt, A. M. Walczak, and T. Mora, TULIP: A transformer-based unsupervised language model for interacting peptides and T cell receptors that generalizes to unseen epitopes, *Proc. Natl. Acad. Sci. USA* **121**, e2316401121 (2024).
- [19] Y. Zhang, Z. Wang, Y. Jiang, D. R. Littler, M. Gerstein, A. W. Purcell, J. Rossjohn, H.-Y. Ou, and J. Song, Epitope-anchored contrastive transfer learning for paired CD8+ T cell receptor-antigen recognition, *Nat. Mach. Intell.* **6**, 1344 (2024).
- [20] Y. Nagano, A. G. Pyo, M. Milighetti, J. Henderson, J. Shawe-Taylor, B. Chain, and A. Tiffeau-Mayer, Contrastive learning of T cell receptor representations, *Cell Syst.* **16**, 101165 (2025).
- [21] M. Nielsen, A. Eugster, M. F. Jensen, M. Goel, A. Tiffeau-Mayer, A. Pelissier, S. Valkiers, M. R. Martínez, B. Meynard-Piganeau, V. Greiff, T. Mora, A. M. Walczak, G. Croce, D. L. Moreno, D. Gfeller, P. Meysman, and J. Barton, Lessons learned from the IMMREP23 TCR-epitope prediction challenge, *Immunoinformatics* **16**, 100045 (2024).
- [22] A. Bellet, A. Habrard, and M. Sebban, A survey on metric learning for feature vectors and structured data, *arXiv:1306.6709*
- [23] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra, Matching networks for one shot learning, *arXiv:1606.04080*.
- [24] S. Henikoff and J. G. Henikoff, Amino acid substitution matrices from protein blocks., *Proc. Natl. Acad. Sci. USA* **89**, 10915 (1992).
- [25] M. Milighetti, J. Shawe-Taylor, and B. Chain, Predicting T cell receptor antigen specificity from structural features derived from homology models of receptor-peptide-major histocompatibility complexes, *Front. Physiol.* **12**, 730908 (2021).
- [26] A. Mayer and C. G. Callan, Measures of epitope binding degeneracy from T cell receptor repertoires, *Proc. Natl. Acad. Sci. USA* **120**, e2213264120 (2023).
- [27] A. Tiffeau-Mayer, Unbiased estimation of sampling variance for Simpson's diversity index, *Phys. Rev. E* **109**, 064411 (2024).
- [28] J. Henderson, Y. Nagano, M. Milighetti, and A. Tiffeau-Mayer, Limits on inferring T cell specificity from partial information, *Proc. Natl. Acad. Sci. USA* **121**, e2408696121 (2024).
- [29] F. Drost, L. Schiefelbein, and B. Schubert, meTCRs - Learning a metric for T-cell receptors, *bioRxiv* (2022), doi:10.1101/2022.10.24.513533.
- [30] M. Pertseva, O. Follonier, D. Scarcella, and S. T. Reddy, TCR clustering by contrastive learning on antigen specificity, *Brief. Bioinform.* **25**, bbae375 (2024).
- [31] A. Murugan, T. Mora, A. M. Walczak, and C. G. Callan, Statistical inference of the generation probability of T-cell receptors from sequence repertoires, *Proc. Natl. Acad. Sci. USA* **109**, 16161 (2012).
- [32] Z. Sethna, G. Isacchini, T. Dupic, T. Mora, A. M. Walczak, and Y. Elhanati, Population variability in the generation and selection of T-cell repertoires, *PLoS Comput. Biol.* **16**, e1008394 (2020).

- [33] Y. Elhanati, A. Murugan, C. G. Callan, T. Mora, and A. M. Walczak, Quantifying selection in immune receptor repertoires, *Proc. Natl. Acad. Sci. USA* **111**, 9875 (2014).
- [34] B. Bravi, A. Di Gioacchino, J. Fernandez-de Cossio-Diaz, A. M. Walczak, T. Mora, S. Cocco, and R. Monasson, A transfer-learning approach to predict antigen immunogenicity and T-cell receptor specificity, *ELife* **12**, e85126 (2023).
- [35] S. F. Edwards and P. W. Anderson, Theory of spin glasses, *J. Phys. F* **5**, 965 (1975).
- [36] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, Supervised Contrastive Learning, *arXiv:2004.11362*.
- [37] M. Ekeberg, C. Lövkvist, Y. Lan, M. Weigt, and E. Aurell, Improved contact prediction in proteins: Using pseudolikelihoods to infer Potts models, *Phys. Rev. E* **87**, 012707 (2013).
- [38] T. Wang and P. Isola, Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere, *arXiv:2005.10242*.
- [39] S. F. Altschul, Amino acid substitution matrices from an information theoretic perspective, *J. Mol. Biol.* **219**, 555 (1991).
- [40] J. Bekker and J. Davis, Learning from positive and unlabeled data: A survey, *Mach. Learn.* **109**, 719 (2020).
- [41] M. Goncharov, D. Bagaev, D. Shcherbinin, I. Zvyagin, D. Bolotin, P. G. Thomas, A. A. Minervina, M. V. Pogorelyy, K. Ladell, J. E. McLaren, D. A. Price, T. H. O. Nguyen, L. C. Rowntree, E. B. Clemens, K. Kedzierska, G. Dolton, C. R. Rius, A. Sewell, J. Samir, F. Luciani, K. V. Zornikova, A. A. Khmelevskaya, S. A. Sheetikov, G. A. Efimov, D. Chudakov, and M. Shugay, VDJdb in the pandemic era: A compendium of T cell receptors specific for SARS-CoV-2, *Nat. Methods* **19**, 1017 (2022).
- [42] W. D. Chronister, A. Crinklaw, S. Mahajan, R. Vita, Z. Koşaloğlu-Yalçın, Z. Yan, J. A. Greenbaum, L. E. Jessen, M. Nielsen, S. Christley, L. G. Cowell, A. Sette, and B. Peters, TCRMatch: predicting T-cell receptor specificity based on sequence similarity to previously characterized receptors, *Front. Immunol.* **12**, 640725 (2021).
- [43] K. Mayer-Blackwell, S. Schattgen, L. Cohen-Lavi, J. C. Crawford, A. Souquette, J. A. Gaevert, T. Hertz, P. G. Thomas, P. Bradley, and A. Fiore-Gartland, TCR meta-clonotypes for biomarker discovery with tcrdist3 enabled identification of public, HLA-restricted clusters of SARS-CoV-2 TCRs, *Elife* **10**, e68605 (2021).
- [44] V. I. Levenshtein, Binary codes capable of correcting deletions, insertions and reversals, *Sov. Phys. Dokl.* **10**, 707 (1966).
- [45] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/14j1-wrh5> for additional figures.
- [46] A. Košmrlj, A. K. Jha, E. S. Huseby, M. Kardar, and A. K. Chakraborty, How the thymus designs antigen-specific and self-tolerant T cell receptor sequences, *Proc. Natl. Acad. Sci. USA* **105**, 16671 (2008).
- [47] J. Leem, S. H. P. de Oliveira, K. Krawczyk, and C. M. Deane, STCRDab: the structural T-cell receptor database, *Nucl. Acids Res.* **46**, D406 (2017).
- [48] H. Mei, Z. H. Liao, Y. Zhou, and S. Z. Li, A new set of amino acid descriptors and its application in peptide QSARs, *Peptide Sci.* **80**, 775 (2005).
- [49] J. Jankauskaitė, B. Jiménez-García, J. Dapkūnas, J. Fernández-Recio, and I. H. Moal, SKEMPI 2.0: An updated benchmark of changes in protein–protein binding energy, kinetics and thermodynamics upon mutation, *Bioinformatics* **35**, 462 (2019).
- [50] A. Postovskaya, K. Vercauteren, P. Meysman, and K. Laukens, tcrBLOSUM: An amino acid substitution matrix for sensitive alignment of distant epitope-specific TCRs, *Brief. Bioinform.* **26**, bbae602 (2024).
- [51] A. S. Perelson and G. F. Oster, Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination, *J. Theor. Biol.* **81**, 645 (1979).
- [52] B. D. Stadinski, K. Shekhar, I. Gómez-Touriño, J. Jung, K. Sasaki, A. K. Sewell, M. Peakman, A. K. Chakraborty, and E. S. Huseby, Hydrophobic CDR3 residues promote the development of self-reactive T cells, *Nat. Immunol.* **17**, 946 (2016).
- [53] V. K. Karnaukhov, D. S. Shcherbinin, A. O. Chugunov, D. M. Chudakov, R. G. Efremov, I. V. Zvyagin, and M. Shugay, Structure-based prediction of T cell receptor recognition of unseen epitopes using TCRen, *Nat. Comput. Sci.* **4**, 510 (2024).
- [54] A. Y. Leary, D. Scott, N. T. Gupta, J. C. Waite, D. Skokos, G. S. Atwal, and P. G. Hawkins, Designing meaningful continuous representations of T cell receptor sequences with deep generative models, *Nat. Commun.* **15**, 4271 (2024).
- [55] H. F. Jones, Z. Molvi, M. G. Klatt, T. Dao, and D. A. Scheinberg, Empirical and rational design of T cell receptor-based immunotherapies, *Front. Immunol.* **11**, 585385 (2021).
- [56] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, Automatic differentiation in PyTorch, in *Conference on Neural Information Processing System, Long Beach, CA, USA* (NIPS, 2017).
- [57] A. Pyo, Inference code for TCR co-specificity prediction, Zenodo repository (2025), [10.5281/zenodo.15707892](https://doi.org/10.5281/zenodo.15707892).